

Please type a plus (+) inside this box -

U.S. PTO



11/03/00

11-06-00

A

PTO/SB/29 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no person is required to respond to a collection of information unless it displays a valid OMB control number.

<b>UTILITY PATENT APPLICATION TRANSMITTAL</b> <i>(Only for new nonprovisional applications under 37 CFR 1.53(b))</i>		Attorney Docket No.	9772-0298-999	Total Pages	
		First Named Inventor or Application Identifier			
		Mark Alexander Najork			
		Express Mail Label No.	EL 451 595 287 US		
<b>APPLICATION ELEMENTS</b> <i>See MPEP chapter 600 concerning utility patent application contents.</i>		<b>ADDRESS TO:</b> Assistant Commissioner for Patents Box Patent Application Washington, DC 20231			
1. <input checked="" type="checkbox"/> Fee Transmittal Form <i>Submit an original, and a duplicate for fee processing</i>		6. <input type="checkbox"/> Microfiche Computer Program <i>(Appendix)</i>			
2. <input checked="" type="checkbox"/> Specification <span style="float: right;">[Total Pages <u>18</u>]</span> <i>(preferred arrangement set forth below)</i> -Descriptive title of the Invention -Cross Reference to Related Applications -Statement Regarding Fed sponsored R&D -Reference to Microfiche Appendix -Background of the Invention -Brief Summary of the Invention -Brief Description of the Drawings <i>(if filed)</i> -Detailed Description of the Invention <i>(including drawings, if filed)</i> -Claim(s) -Abstract of the Disclosure		7. <input type="checkbox"/> Nucleotide and/or Amino Acid Sequence Submission <i>(if applicable, all necessary)</i> a. <input type="checkbox"/> Computer Readable Copy b. <input type="checkbox"/> Paper Copy <i>(identical to computer copy)</i> c. <input type="checkbox"/> Statement verifying identity of above copies			
3. <input checked="" type="checkbox"/> Drawing(s) <i>(35 USC 113)</i> <span style="float: right;">[Total Sheets <u>9</u>]</span>		<b>ACCOMPANYING APPLICATION PARTS</b> 8. <input checked="" type="checkbox"/> Assignment Papers <i>(cover sheet &amp; document(s))</i> 9. <input checked="" type="checkbox"/> 37 CFR 3.73(b) Statement <input checked="" type="checkbox"/> Power of Attorney <i>(when there is an assignee)</i> 10. <input type="checkbox"/> English Translation Document <i>(if applicable)</i> 11. <input checked="" type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input checked="" type="checkbox"/> Copies of IDS Citations 12. <input type="checkbox"/> Preliminary Amendment 13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) <i>(Should be specifically itemized)</i> 14. <input type="checkbox"/> Small Entity <input type="checkbox"/> Statement filed in prior application, Statement(s) Status still proper and desired 15. <input type="checkbox"/> Certified Copy of Priority Document(s) <i>(if foreign priority is claimed)</i> 16. <input type="checkbox"/> Other:			
4. <input checked="" type="checkbox"/> Oath or Declaration <span style="float: right;">[Total Sheets <u>2</u>]</span> <input checked="" type="checkbox"/> Newly executed <i>(original or copy)</i> <input type="checkbox"/> Copy from a prior application (37 CFR 1.63(d)) <i>(for continuation/divisional with Box 17 completed)</i> <b>[Note Box 5 below]</b> i. <input type="checkbox"/> <b>DELETION OF INVENTORS(S)</b> Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33 (b).					
5. <input type="checkbox"/> Incorporation By Reference <i>(useable if Box 4b is checked)</i> The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.					
17. If a <b>CONTINUING APPLICATION</b> , check appropriate box and supply the requisite information: <input type="checkbox"/> Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No: filed .					
<b>18. CORRESPONDENCE ADDRESS</b>					
<input checked="" type="checkbox"/> Customer Number or Bar Code Label		24341 <i>(Insert Customer No. or Attach bar code label here)</i>		or <input type="checkbox"/> Correspondence address below	
NAME					
ADDRESS					
CITY		STATE		ZIP CODE	
COUNTRY		TELEPHONE		FAX	

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

**PENNIE & EDMONDS LLP**

COUNSELLORS AT LAW  
 1155 Avenue of the Americas  
 New York, N.Y. 10036-2711  
 (212) 790-9090

ATTORNEY DOCKET NO. 9772-0298-999Date: November 3, 2000

Assistant Commissioner for Patents  
 Box PATENT APPLICATION  
 Washington, D.C. 20231

Sir:

The following utility patent application is enclosed for filing:

Applicant(s): Mark Alexander Najork

Executed on: November 2, 2000

Title of Invention: SYSTEM AND METHOD FOR DISTRIBUTED WEB CRAWLING

**PATENT APPLICATION FEE VALUE**

TYPE	NO. FILED	LESS	EXTRA	EXTRA RATE	FEE
Total Claims	14	-20	0	\$18.00 each	\$ 0.00
Independent	4	-3	1	\$80.00 each	\$ 80.00
Minimum Fee					\$ 710.00
Multiple Dependency Fee If Applicable (\$270.00)					\$ 0.00
<b>Total</b>					\$ 790.00
50% Reduction for Independent Inventor, Nonprofit Organization or Small Business Concern (a verified statement as to the applicant's status is attached)					- \$ 0.00
<b>Total Filing Fee</b>					\$ 790.00

- ☐ Priority of application no. filed on in is claimed under 35 U.S.C. § 119.
- ☐ The certified copy of the priority application has been filed in application no. filed
- ☐ Amend the specification by inserting before the first line the following sentence: This is a continuation-in-part of application no. filed.

Please charge the required fee to Pennie & Edmonds LLP Deposit Account No. 16-1150. A copy of this sheet is enclosed.

Respectfully submitted,

Gary S. Williams  
 PENNIE & EDMONDS LLP

31,066  
 (Reg. No.)

Enclosure

This form is not for use with continuation, divisional, re-issue, design or plant patent applications.

## SYSTEM AND METHOD FOR DISTRIBUTED WEB CRAWLING

The present invention relates to a system and method for distributed web crawling and, more particularly, to a web crawling system that uses multiple web crawlers to efficiently process addresses to be downloaded.

### BACKGROUND OF THE INVENTION

Documents on interconnected computer networks are typically stored on numerous host computers that are connected over the networks. For example, so-called "web pages" may be stored on the global computer network known as the Internet, which includes the world wide web. Web pages can also be stored on Intranets, which are typically private networks maintained by corporations, government entities, and other groups. Each web page, whether on the world wide web or an Intranet, has a distinct address called its uniform resource locator (URL), which at least in part identifies the location or host computer of the web page. Many of the documents on Intranets and the world wide web are written in standard document description languages (e.g., HTML, XML). These languages allow an author of a document to create hypertext links to other documents. Hypertext links allow a reader of a web page to access other web pages by clicking on links to the other pages. These links are typically highlighted in the original web page. A web page containing hypertext links to other web pages generally refers to those pages by their URL's. A URL may be referred to more generally as a data set address, which corresponds to a web page, or data set. Links in a web page may refer to web pages that are stored in the same or different host computers.

A web crawler is a program that automatically finds and downloads documents from host computers in an Intranet or the world wide web. A computer with a web crawler installed on it may also be referred to as a web crawler. When a web crawler is given a set of starting URL's, the web crawler downloads the corresponding documents. The web crawler then extracts any URL's contained in those downloaded documents. Before the web crawler

downloads the documents associated with the newly discovered URL's, the web crawler needs to find out whether these documents have already been downloaded. If the documents associated with the newly discovered URL's have not been downloaded, the web crawler downloads the documents and extracts any URL's contained in them. This process repeats indefinitely or until a predetermined stop condition occurs. This process is demanding for a single web crawler due to the large number of URL's to download and process. As of 1999 there were approximately 800 million web pages on the world wide web and the number is continuously growing. Even Intranets can store millions of web pages.

## SUMMARY OF THE INVENTION

The present invention provides for the efficient downloading of data set addresses from among a plurality of host computers, using a plurality of web crawlers. Each web crawler identifies URL's in data sets downloaded by that web crawler, and identifies the host computer identifier within each such URL. The host computer identifier for each URL is mapped to the web crawler identifier of one of the web crawlers. If the URL is mapped to the web crawler identifier of a different web crawler, the URL is sent to that web crawler for processing, and otherwise the URL is processed by the web crawler that identified the URL.

Each web crawler sends URL's to the other web crawlers for processing, and each web crawler receives URL's from the other web crawlers for processing. Each web crawler processes only the URL's assigned to it, which are the URL's whose host identifier is mapped to the web crawler identifier for that web crawler. Each web crawler filters the URL's assigned to it by comparing them against a database of URL's already known by the web crawler. If a URL is not already known to the web crawler, the data set corresponding to the URL is scheduled for downloading; otherwise, the URL is ignored.

## BRIEF DESCRIPTION OF THE DRAWINGS

Additional objects and features of the invention will be more readily apparent from the following detailed description and appended claims when taken in conjunction with the drawings, in which:

Fig. 1 is a block diagram of a web crawler system in accordance with an embodiment of the invention.

Fig. 2 is a block diagram of a web crawler in accordance with an embodiment of the present invention.

Fig. 3 is a flow chart of a main web crawling procedure executed by each of a plurality of threads in an exemplary embodiment of the invention.

Fig. 4 is a flow chart of an address distribution procedure used in an exemplary embodiment of the invention.

Fig. 5 is a flow chart of an address filter procedure used in an exemplary embodiment of the invention.

Fig. 6 is a block diagram of the data flow of an embodiment of the address filter procedure.

Fig. 7 is a block diagram of the data structures used to store the addresses known to the web crawler.

Fig. 8 is a block diagram of the data flow of an alternate embodiment of the address filter procedure.

Fig. 9 is a block diagram of the data flow of an alternate embodiment of the address filter procedure.

## DESCRIPTION OF THE PREFERRED EMBODIMENTS

Referring to Fig. 1, there is shown an embodiment of a distributed computer system 100. The distributed computer system 100 includes a web crawler system 118 connected to a network 120 through a network interconnection 119. The network 120 may be a global

communication network, such as the Internet, or a private network, sometimes called an Intranet. Examples of network interconnection 119 include switches, routers, etc.

The network 120 includes web servers 121 and may optionally include a service known as a domain name system 122. It may also optionally include a web page indexing system 123. The web servers 121 store web pages. The domain name system 122 provides the mapping between Internet Protocol (IP) addresses and host names. Each site participating in the domain name system 122 maintains its own database of information and runs a server program that other systems across the Intranet or Internet can query. The domain name system provides the protocol that allows clients and servers to communicate with each other. Any application may look up the IP address (or addresses) corresponding to a given host name or the host name corresponding to a given IP address in the domain name system 122. An application accesses the domain name system 122 through a resolver. The resolver contacts one or more name servers to perform a mapping of a host name to the corresponding IP address, or vice versa. A given host name may be associated with more than one IP address because an Intranet or Internet host may have multiple interfaces, with each interface of the host having a unique IP address.

The web page indexing system 123 includes an index of words used on the world wide web, or the Intranet, and addresses of the web pages that use each word. Such indexing systems are maintained by various search engines, such as the AltaVista search engine. The domain name system 122 and the web page indexing system 123 may be accessed by the web crawlers 116 in the process of downloading web pages from the world wide web.

The web crawler system 118 is made up of  $n$  web crawlers, where  $n$  is an integer larger than one. A typical web crawler system 118 may contain, for example, four web crawlers. The web crawlers 116 in the web crawler system 118 communicate via a communications network such as a local area network, or the Internet. In addition, there may be a separate computer (not shown) for controlling the communication between the web crawlers in the web crawler system 118.

An exemplary embodiment of a web crawler 116 is shown in Figure 2. The web crawler 116 includes one or more CPUs 101, a communications interface or network connection 102, memory 104. The web crawler may optionally include a user interface 103, or the web crawler may be accessed by a user interface or other facility remotely located on another computer. In some implementations two network connections or communication

interfaces may be needed, one for downloading web documents, and another for communicating with the other web crawlers. The memory 104 includes:

- an operating system 105;
- an Intranet/Internet access procedure 106 for fetching web pages as well as communicating with the domain name system 122 (Fig. 1);
- a frontier 140 for storing URL's or representations of URL's that will be downloaded by the web crawler 116;
- threads 108 for downloading web pages from the servers 121, and processing the downloaded web pages;
- a main web crawler procedure 109 executed by each of the threads 108;
- a URL address distribution procedure 110 executed by each of the threads 108 to identify the URL's in a downloaded web page, and to determine which web crawler is associated with each URL;
- a web crawler lookup table 111 for storing the addresses corresponding to the web crawler identifiers of the other web crawlers in the system;
- a URL address filter procedure 112 executed by each of the threads 108 to filter the URL's assigned to this web crawler for processing, and to determine which URL's are new and therefore should be scheduled for downloading; and
- data structures 113 for storing all of the known addresses associated with the web crawler.

The data structures 113 for storing the known data set addresses may include:

- a buffer 172 for storing the addresses (and/or fingerprints of addresses) identified as potentially being addresses unknown to the web crawler;
- an optional cache 174, for storing frequently encountered addresses (used in an embodiment described with reference to Fig. 8); and
- a disk file (or set of disk files) 176 for storing addresses (and/or fingerprints of addresses) known to the web crawler.

While these are the data structures used in a preferred embodiment to store known data set addresses, and/or the fingerprints of those addresses, other data structures may be used to store this information in other embodiments.

## Main Web Crawler Procedure

In the exemplary embodiment, each web crawler uses multiple threads to download and process documents. The web crawler 116 is given a set of initial URL's and begins  
5 downloading documents using those URL's. Various data structures may be used to keep track of which documents (web pages) the threads should download and process, but those particular data structures are not the subject of the present invention. Rather, the present invention concerns the methodology and data structures used to determine the URL's to be processed by each web crawler in the web crawler system.

10 Each thread executes a main web crawler procedure 109, which will be now described with reference to the flow chart shown in Fig. 3 and the block diagram of Fig. 2. The web crawler thread determines the data set address of the next document, or data set, to be downloaded, typically by retrieving the data set address from the frontier 140 (step 150). The frontier 140 is a queue that stores the data set addresses (URL's) to be downloaded. The  
15 thread then downloads the document corresponding to the URL (step 151), and processes the document. The processing may include indexing the words in the document so as to make the document accessible via a search engine. However, the only processing of the document that is relevant to the present discussion is that the main procedure identifies URL's in the downloaded document (step 152) that are candidates for downloading and processing.  
20 Typically, these URL's are found in hypertext links in the document being processed.

Each identified URL is passed to the address distribution procedure in step 153 to determine which web crawler in the web crawler system should process the URL. If the address distribution procedure determines that the identified URL should not be processed by the web crawler running the procedure, the URL is sent to the appropriate web crawler for  
25 processing. In the meantime, the main procedure continues to identify URL's and pass URL's to the address distribution procedure until all URL's from the downloaded page have been processed (steps 154 and 155). When all URL's have been passed to the address distribution procedure, the next address is retrieved from the frontier 140 (step 150), and then the procedure described above repeats. The main web crawler procedure continues to run as  
30 long as there are URL's in the frontier for it to process, or until a predetermined stop condition occurs.



## Address Distribution Procedure

An exemplary embodiment of the address distribution procedure invoked by the main web crawler procedure is described next with reference to the flow chart shown in Fig. 4 and the block diagram in Fig. 2. It is preferable to assign URL's to each web crawler by host name. Each web crawler is then responsible for processing only the addresses associated with its assigned set of hosts. Assigning all the URL's from each host to one particular web crawler makes it more practical to enforce "politeness" policies, which dictate restrictions on how often the web crawler system will download pages from any one host, and it avoids having to replicate per-host information across multiple web crawlers.

The first step of the URL distribution procedure is to convert the part of the specified URL U that identifies the web page's host computer, called the host computer identifier, into a numeric representation H (step 160). In one embodiment, the "host computer identifier" is a symbolic name, such as "www.compaq.com," while in another embodiment, the "host computer identifier" is an IP address, such as 161.114.19.252. Which of these two embodiments is used determines the interpretation of "politeness." For instance, if the system distributes URL's based on symbolic host names, then it is possible that the web crawler system will distribute two different host names that happen to resolve to the same IP address to two different web crawlers, in which case those two web crawlers may send requests in parallel to the machine at that IP address. If the system distributes URL's based on the host IP address, only one web crawler will send requests to the host computer at that IP address.

In the preferred embodiment, an intermediate value V is first produced by converting the host computer identifier into a fixed length numeric representation, or "hash value." The hash value, also herein called the fingerprint, may be generated by applying a predefined hash function to the specified URL U's host computer identifier. The only requirement for the hash function is that it should spread the URL's evenly over the available hash buckets. Many suitable hash functions are known to those skilled in the art, including checksum functions and fingerprint functions, and thus are not described here.

After the host identifier of the specified URL U has been converted into the intermediate value V, the web crawler generates the representation of the host computer identifier by computing a function of the host computer identifier whose result is an integer value that is a member of a set of n predefined distinct values, where n is the number of web

crawlers in the web crawler system. For instance, the representation of the host computer identifier may be generated by computing  $V$  modulo  $n$ . The host computer representation  $H$  which is the result of computing  $V$  modulo  $n$  will be an integer between 0 and  $n-1$ . In alternate embodiments, other methods may be used to map the host identifier into a numerical representation  $H$ . One such alternative method is to divide the numeric range of all possible intermediate values  $V$  into  $n$  sub-ranges, where the range of values is preferably divided so as to evenly spread the possible values of  $V$  over the  $n$  sub-ranges. Each value  $V$  is then mapped to an integer between 0 and  $n-1$  by determining which of the  $n$  sub-ranges it falls in. For instance, if the hash function used to generate  $V$  returns a value between 0 and  $X-1$ , the web crawler identifier may be computed as follows:

```
if       $0 \leq V < (1/n)X$       then return 0
if  $(1/n)X \leq V < (2/n)X$     then return 1
if  $(2/n)X \leq V < (3/n)X$     then return 2
...
if  $((n-1)/n)X \leq V < X$     then return  $n-1$ 
```

After the host computer representation  $H$  has been generated, the address distribution procedure determines the web crawler identifier corresponding to  $H$ . In the preferred embodiment, each web crawler in the web crawler system is assigned a web crawler identifier  $W$ . The web crawler identifier may be an integer between zero and  $n-1$ , where  $n$  is the number of web crawlers in the web crawler system. Therefore, in this embodiment, the host computer representation  $H$  is equal to its corresponding web crawler identifier  $W$ . Other embodiments may include a procedure to map the host computer representation  $H$  to the web crawler identifier  $W$  to ensure an even distribution of URL's across the plurality of web crawlers. If the web crawler identifier  $W$  is assigned to the current web crawler (i.e., the web crawler running the address distribution procedure), then the URL  $U$  is passed to the address filter procedure, which is described below (step 166). If the web crawler identifier  $W$  is not assigned to the current web crawler, the URL  $U$  is sent to the destination web crawler corresponding to the web crawler identifier  $W$  (step 167).

Other methods of generating a host computer representation and mapping it to the web crawler identifier such that there is an even spreading of host names (host computer identifiers) over the web crawlers in the web crawler system may be apparent to those skilled in the art.

In order to send the URL U to the destination web crawler, the sending web crawler transmits the URL U to the destination web crawler at an address specified by the web crawler lookup table 111. The web crawler lookup table 111 may be part of a configuration file used to determine the configuration of the web crawler during initialization of the system.

5 In a preferred embodiment, during initialization of the web crawler system each web crawler opens a connection to every other web crawler in the system, and thereafter transmits URL's to the other web crawlers over the open connections. The "open connections" are operating system or other software constructs that generate message packets that are sent to the destination web crawlers via the network connection 102. Each such message packet  
10 contains the information (e.g., a URL) being sent and has a destination address equal to the IP address of the web crawler to which the URL is being sent.

#### Address Filter Procedure

15 The address filter procedure is described with reference to the flow diagram in Fig. 5 and the block diagram in Fig. 2. The address filter procedure receives URL's sent from other web crawlers and from the address distribution procedure as described above (step 167). The address filter procedure determines if the URL U is a duplicate URL that has already been downloaded by the web crawler or has already been scheduled for downloading (step 168).  
20 This is preferably done by comparing the URL to a list of all stored URL's that are known to the web crawler. The list of known URL's is stored in a set of data structures 113, described in more detail below, designed to facilitate efficient lookup. If the URL is a duplicate (168-Yes), it is discarded (step 170). If the URL is not a duplicate (168-No), it is added to the frontier 140 for downloading by the web crawler and to the list of the addresses known to the  
25 web crawler (step 169). The data structures for the frontier and for the list of addresses known to the web crawler are described below.

The diagram in Fig. 6 shows an exemplary embodiment of the data structures used to determine whether the specified URL is a duplicate. Preferably, a fingerprint N is generated to represent the entire specified URL, not just the host computer identifier. Next, a lookup is  
30 performed to see if N is already stored in a buffer B (172), which stores the fingerprints of recently added URL's, as well as the corresponding URL's. If N is already stored in buffer B 172, it has already been scheduled for downloading, and therefore no further processing of the URL U is needed (i.e., the new copy of the URL U is discarded). If N is not stored in

buffer B 172, it may be new (although it may already be stored in the fingerprint disk file 176), and therefore N and the corresponding URL U are added to buffer B 172. Other URL's are processed until buffer B 172 is full.

When the buffer B 172 becomes full, it must be merged with the fingerprint disk file 176. The fingerprints in buffer B 172 are first sorted so that the fingerprints are in order. After ordering, the buffer B 172 is then merged with the fingerprint disk file 176, where only fingerprints corresponding to new URL's are added to the disk file 176. During the merge, the fingerprint disk file 176 is searched to determine whether each fingerprint from buffer B is located in the disk file. Each fingerprint located in the disk file is discarded. Each fingerprint not located in the disk file is added to the disk file 176, and the corresponding URL is added to the frontier for downloading of the corresponding web page document. When the merge process is completed, the contents of buffer B 172 are deleted. The process of merging buffer B 172 into the disk file 176 is a relatively time-consuming operation. Therefore buffer B 172 is typically made fairly large so as to minimize the frequency of such merge operations. In this embodiment, buffer B 172 and the disk file 176 are the data structures 113 (Fig. 2) used to store known addresses.

During the merge process, which is an ordered merge, each fingerprint from buffer B not found in the disk file must be inserted in the fingerprint disk file 176 in the proper location, as illustrated in Figure 7, so that the disk file 176 remains ordered. This requires the disk file to be completely re-written. To avoid this lengthy rewrite process, in a preferred embodiment, the fingerprint disk file may be sparsely-filled, using open addressing. For this embodiment, the fingerprint disk file represents a hash table, with a substantial proportion of the table, for example 50% or 75%, being empty entries or "holes."

In this embodiment, in order to determine whether a fingerprint from buffer B is in the disk file, the hash of the fingerprint is computed. In one embodiment, a prefix of the fingerprint is used for the hash value. The hash value is the starting position for searching through the fingerprint disk file. The disk file is searched sequentially, starting at the starting position, for either a match or a hole. If a hole is found, the fingerprint from buffer B is stored in that hole; if a match is found, it is discarded. Thus, there is only one write to the disk file for each fingerprint not already present in the disk file, and the size of the disk file is not a factor in the merge time. When the disk file becomes too full - for example, when only 25% of the slots in the disk file are holes - the file must be completely rewritten into a new, larger file. For example, the new file may be doubled in size, in which case the amortized

cost of maintaining the file is constant per fingerprint in the hash table. It will be appreciated that the use of a sparsely-filled disk file drastically reduces the disk re-writing required during a merge.

In one embodiment, the disk file may be divided into sparse sub-files, with open-addressing used for each sub-file. An index may be used to identify the range of fingerprint hash values located in each sub-file, or an additional hash table may be used to map fingerprints to the various sub-files. When a sub-file becomes too full, it may be re-written into a new, larger file, but the entire disk file need not be re-written.

The disk file (or each sub-file) is ordered, so that all fingerprints from the same host will be likely stored in either the same sub-file, or at contiguous memory locations within the disk file. Because URL's located on a given web page often have the same host, this efficient addressing scheme can significantly improve the speed of the merge process. When a new fingerprint is to be added to the disk file during the merge, often the buffer used for the merge operation will already contain the disk page on which that fingerprint will be stored. This will occur when the previous fingerprint added to the disk file was from the same host as the new fingerprint, and consequently would likely have been stored on the same disk page as the new fingerprint.

In an alternative embodiment, illustrated in Fig. 8, in addition to buffer B 172, a cache C 174 of "popular" URL's is included. Cache C, in this embodiment, is one of the data structures 113 (Fig. 2) for storing known addresses. The fingerprint for each URL from the address distribution procedure or from another web crawler 141 is first compared to fingerprints in cache C 174, to see whether the URL is a "popular" URL that has already been downloaded (or scheduled for downloading) and therefore need not be processed any further. If the fingerprint is not in cache C 174, the address filter then goes on to compare the fingerprint to the fingerprints in buffer B 172, as discussed previously. Use of the cache C 174 significantly reduces the rate at which the buffer B 172 is filled with fingerprints, which reduces the frequency of merges with the disk file and thus improves the efficiency of the address filter.

The cache C is updated according to an update policy. When, according to the update policy, a new fingerprint must be added to cache C, and cache C 174 is full, a fingerprint is evicted from cache C 174. The fingerprint to be evicted at any one time is determined using an appropriate eviction policy, such as a least recently used eviction policy, a round robin eviction policy or a clock replacement eviction policy.

A new fingerprint is added to cache C 174 according to an update policy. The following are examples of update policies. A fingerprint corresponding to a URL may be added to cache C 174, following the comparison of the fingerprint to the contents of cache C 174, whenever that fingerprint is not already contained in cache C 174. Thus, under this  
5 update policy, cache C 174 is simply a cache of fingerprints for recently-identified URL's. Alternatively, a fingerprint may be added to cache C 174 following the comparison of a fingerprint to the contents of buffer B 172 whenever that fingerprint is not contained in cache C 174, but is contained in buffer B 172. In this case, cache C 174 is a cache of fingerprints for URL's which have been identified more than once recently. Or, a fingerprint may be  
10 added to cache C whenever that fingerprint is not contained in cache C 174 and is not contained in buffer B 172. Cache C 174 in this case is a cache of fingerprints for URL's which have been identified only once recently. Finally, a fingerprint may be added to cache C 174 during the merge when the fingerprint is found to have been previously stored in the fingerprint disk file 176. For this policy, cache C 174 is a cache of fingerprints for URL's  
15 which have been identified more than once during the entire web crawl.

Another alternative embodiment is illustrated in Fig. 9. When buffer B 172 is full, it must be merged 142 with the fingerprint disk file 176, and cannot accept any newly-identified URL's. Hence, when using the embodiment shown in Fig. 7 or Fig. 8, no new addresses can be accepted for filtering while the time-consuming merge process is in progress. To avoid  
20 delaying the address filtering procedure during the merge process, in the embodiment shown in Fig. 9 the address filter uses a second buffer B' 180 in addition to the buffer B 172. In this embodiment, buffer B' 180 is one of the data structures 113 (Fig. 2) used to store known addresses. While buffer B 172 is being merged with the disk file 176, buffer B' 180 takes over the role of buffer B 172. Fingerprints for newly-identified URL's are compared with the  
25 contents of buffer B' 180 and then stored in buffer B' 180 if it is determined that the fingerprints are not already stored in buffer B' 180. When buffer B' is full, the two buffers change roles again.

The cache C 174 of the embodiment in Fig. 8 can also be used in the embodiment shown in Fig. 9. The addition of cache C makes use of the buffer B 172/180 more efficient,  
30 while the use of two buffer B's 172, 180 greatly reduces the "cost" of the buffer merge operations in terms of the web crawler's ability to process discovered URL's that may or may not be new URL's.

In yet another alternate embodiment, the list of known addresses is maintained in a hash table in main memory (i.e., not on disk). This embodiment is useful primarily in systems used to crawl Intranets, but can also be used in web crawler systems where the cost of having a very large main memory (e.g., several gigabytes of random access memory) is economically worthwhile. The advantage of keeping the list of known addresses entirely in main memory is speed of operation, while the disadvantage is cost.

The present invention can be implemented as a computer program product that includes a computer program mechanism embedded in a computer readable storage medium. For instance, the computer program product could contain the program modules shown in Fig. 2. These program modules may be stored on a CD-ROM, magnetic disk storage product, or any other computer readable data or program storage product. The software modules in the computer program product may also be distributed electronically, via the Internet or otherwise, by transmission of a computer data signal (in which the software modules are embedded) on a carrier wave.

While the present invention has been described with reference to a few specific embodiments, the description is illustrative of the invention and is not to be construed as limiting the invention. Various modifications may occur to those skilled in the art without departing from the true spirit and scope of the invention as defined by the appended claims.

WHAT IS CLAIMED IS:

1. A method of downloading data sets by a plurality of web crawlers from among a plurality of host computers, comprising the steps of:
  - assigning a web crawler identifier to each one of the plurality of web crawlers;
  - for each respective web crawler:
    - downloading at least one data set that includes addresses of one or more referred data sets;
    - identifying the addresses of the one or more referred data sets, wherein each identified address includes a host computer identifier;
    - for each identified address:
      - generating a representation of the host computer identifier;
      - determining a web crawler identifier to which the representation corresponds; and
      - when the determined web crawler identifier is not assigned to the respective web crawler, sending the identified address to the web crawler to which the determined web crawler identifier is assigned.
2. The method of claim 1, wherein
  - the plurality of web crawlers consists of  $n$  web crawlers; and
  - generating the representation includes computing a function of the host computer identifier to generate an integer value that is a member of a set of  $n$  predefined distinct values.
3. The method of claim 1, wherein
  - the plurality of web crawlers consists of  $n$  web crawlers; and
  - generating the representation includes computing a hash function of the host computer identifier to generate an intermediate value  $V$ , and computing  $V$  modulo  $n$ .
4. The method of claim 1, wherein the sending step includes:
  - determining a web crawler address for the web crawler to which the determined web crawler identifier is assigned;
  - transmitting the identified data set address to the destination web crawler at the determined web crawler address.



1 5. A method of downloading data sets by a plurality of web crawlers from among a  
2 plurality of host computers, comprising the steps of:  
3 for each respective web crawler:  
4 receiving addresses of one or more data sets from each of the plurality of web  
5 crawlers other than the respective web crawler;  
6 for each received address:  
7 determining if the address has been previously stored; and  
8 if this determination is negative, storing the address.

1 6. A web crawler system for downloading data set addresses from among a plurality of  
2 host computers, comprising:  
3 a plurality of web crawlers, wherein each web crawler has been assigned a web  
4 crawler identifier;  
5 for each respective web crawler:  
6 a main web crawler module for downloading and processing data sets stored  
7 on a plurality of host computers, the main web crawler module identifying addresses of the  
8 one or more referred data sets in the downloaded data sets, wherein each identified address  
9 includes a host computer identifier; and  
10 an address distribution module for processing the identified addresses, the  
11 address distribution module including instructions for:  
12 generating a representation of the host computer identifier, wherein the  
13 representation corresponds to one of the web crawler identifiers;  
14 determining a web crawler identifier to which the representation  
15 corresponds; and  
16 when the determined web crawler identifier is not assigned to the  
17 respective web crawler, sending the identified address to a destination web crawler  
18 comprising the web crawler to which the determined web crawler identifier is assigned.

1 7. The web crawler system of claim 6 wherein  
2 the plurality of web crawlers consists of n web crawlers; and

the address distribution module's instructions for generating the representation includes instructions for computing a hash function of the host computer identifier to generate an intermediate value V, and computing V modulo n.

8. The web crawler system of claim 6, further comprising:  
for each respective web crawler, a web crawler interface for transmitting the identified address to the destination web crawler and for receiving identified addresses from each of the plurality of web crawlers other than the respective web crawler.

9. The web crawler system of claim 6, further comprising:  
for each respective web crawler, a lookup table storing for each of the plurality of web crawler identifiers a corresponding web crawler address, said lookup table for use by the address distribution module in determining a web crawler address to which to send the identified data set address.

10. A computer program product for use in conjunction with a web crawler system wherein each web crawler is assigned a web crawler identifier, the computer program product comprising a computer readable storage medium and a computer program mechanism embedded therein, the computer program mechanism comprising:

a main web crawler module for downloading and processing data sets stored on a plurality of host computers, the main web crawler module identifying addresses of the one or more referred data sets in the downloaded data sets, wherein each identified address includes a host computer identifier; and

an address distribution module for processing the identified addresses, the address distribution module including instructions for:

generating a representation of the host computer identifier, wherein the representation corresponds to one of the web crawler identifiers;

determining a web crawler identifier to which the representation corresponds;

and

when the determined web crawler identifier is not assigned to the respective web crawler, sending the identified address to a destination web crawler comprising the web crawler to which the determined web crawler identifier is assigned.

1 11. The computer program product of claim 10, wherein:  
2 the web crawler system consists of n web crawlers; and  
3 the address distribution module's instructions for generating the representation  
4 includes instructions for computing a function of the host computer identifier to generate an  
5 integer value that is a member of a set of n predefined distinct values.

1 12. The computer program product of claim 10, wherein:  
2 the web crawler system consists of n web crawlers; and  
3 the address distribution module's instructions for generating the representation  
4 includes instructions for computing a hash function of the host computer identifier to  
5 generate an intermediate value V, and computing V modulo n.

1 13. The computer program product of claim 10, further comprising:  
2 a web crawler interface for transmitting the identified address to the destination web  
3 crawler and for receiving identified addresses from each of the plurality of web crawlers other  
4 than the respective web crawler.

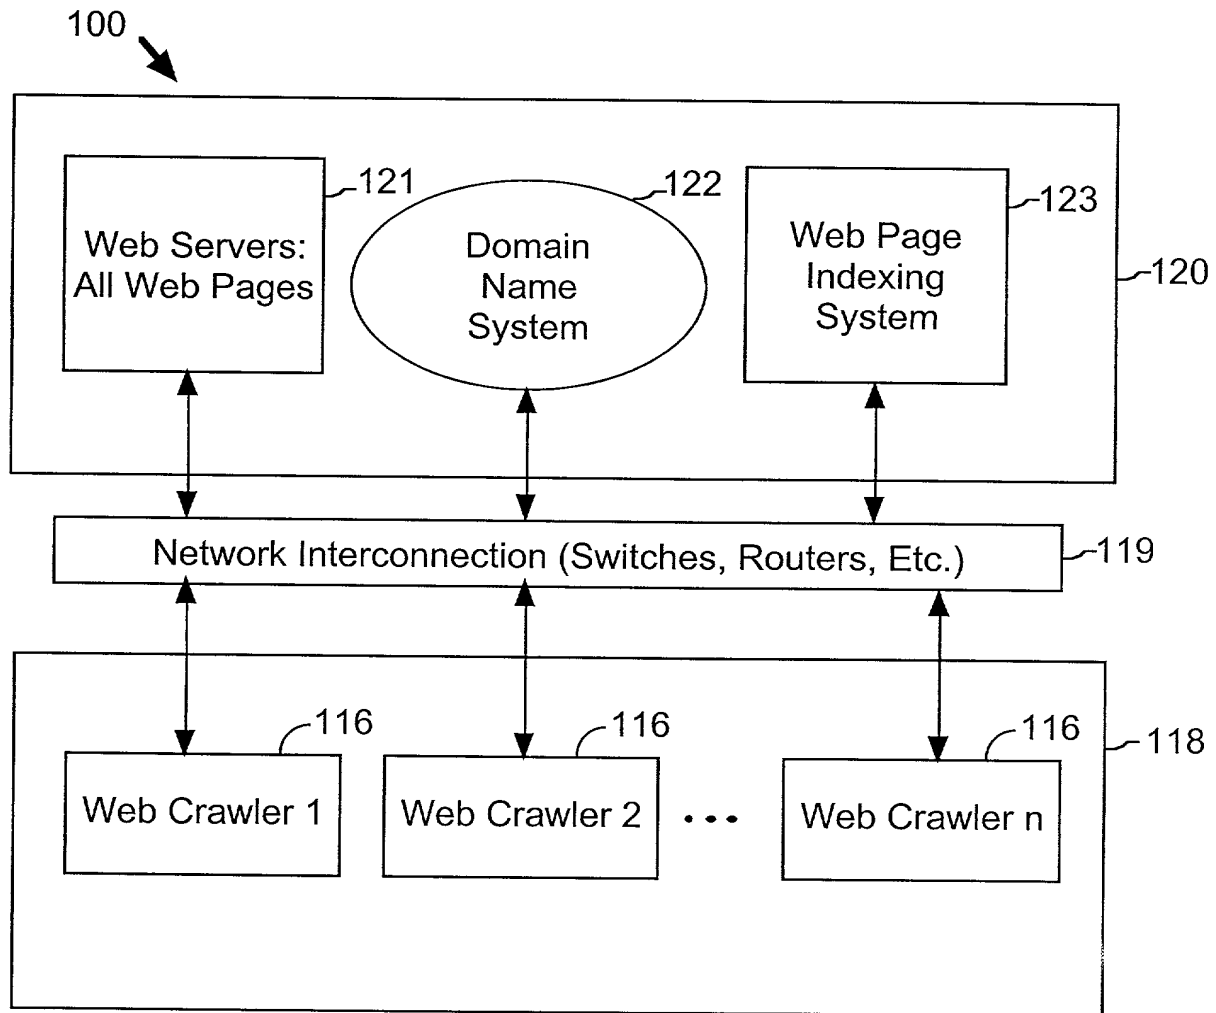
1 14. The computer program product of claim 10, further comprising:  
2 a lookup table storing for each of the plurality of web crawler identifiers a  
3 corresponding web crawler address, said lookup table for use by the address distribution  
4 module in determining a web crawler address to which to send the identified data set address.

# SYSTEM AND METHOD FOR DISTRIBUTED WEB CRAWLING

## ABSTRACT OF THE DISCLOSURE

5 The present invention provides for the efficient downloading of data set addresses from among a plurality of host computers, using a plurality of web crawlers. Each web crawler identifies URL's in data sets downloaded by that web crawler, and identifies the host computer identifier within each such URL. The host computer identifier for each URL is mapped to the web crawler identifier of one of the web crawlers. If the URL is mapped to the web crawler identifier of a different web crawler, the URL is sent to that web crawler for processing, and otherwise the URL is processed by the web crawler that identified the URL. Each web crawler sends URL's to the other web crawlers for processing, and each web crawler receives URL's from the other web crawlers for processing. In a preferred embodiment, each web crawler processes only the URL's assigned to it, which are the URL's whose host identifier is mapped to the web crawler identifier for that web crawler. Each web crawler filters the URL's assigned to it by comparing them against a database of URL's already known by the web crawler and removing the already known URL's. If a URL is not already known to the web crawler, the data set corresponding to the URL is scheduled for downloading.

Compaq/DEC 9772-0298-999

**FIG. 1**





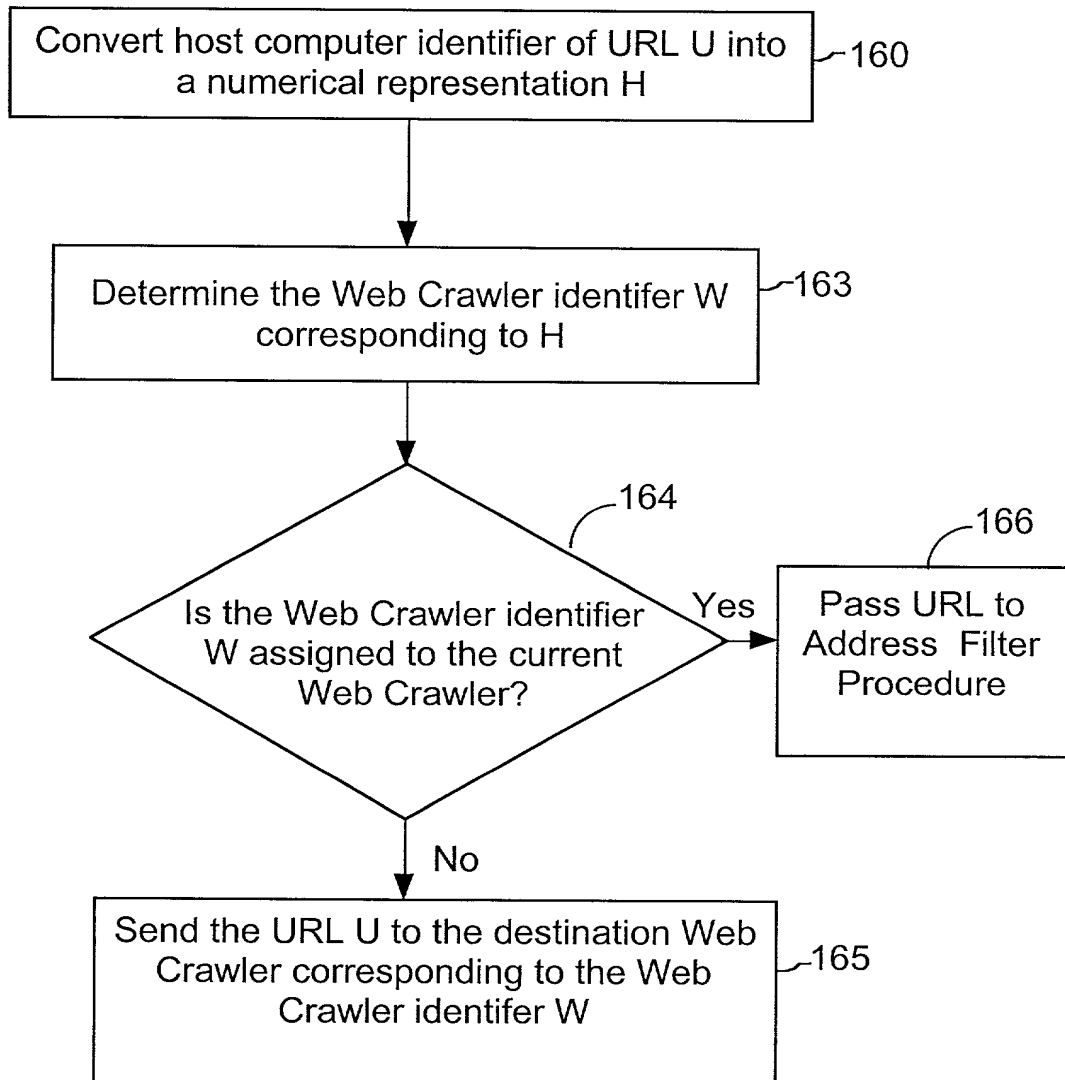
FIG. 2

Main Procedure, performed by each thread  
109



**FIG. 3**

Address Distribution Procedure,  
Step 153



**FIG. 4**



## Address Filter Procedure

112

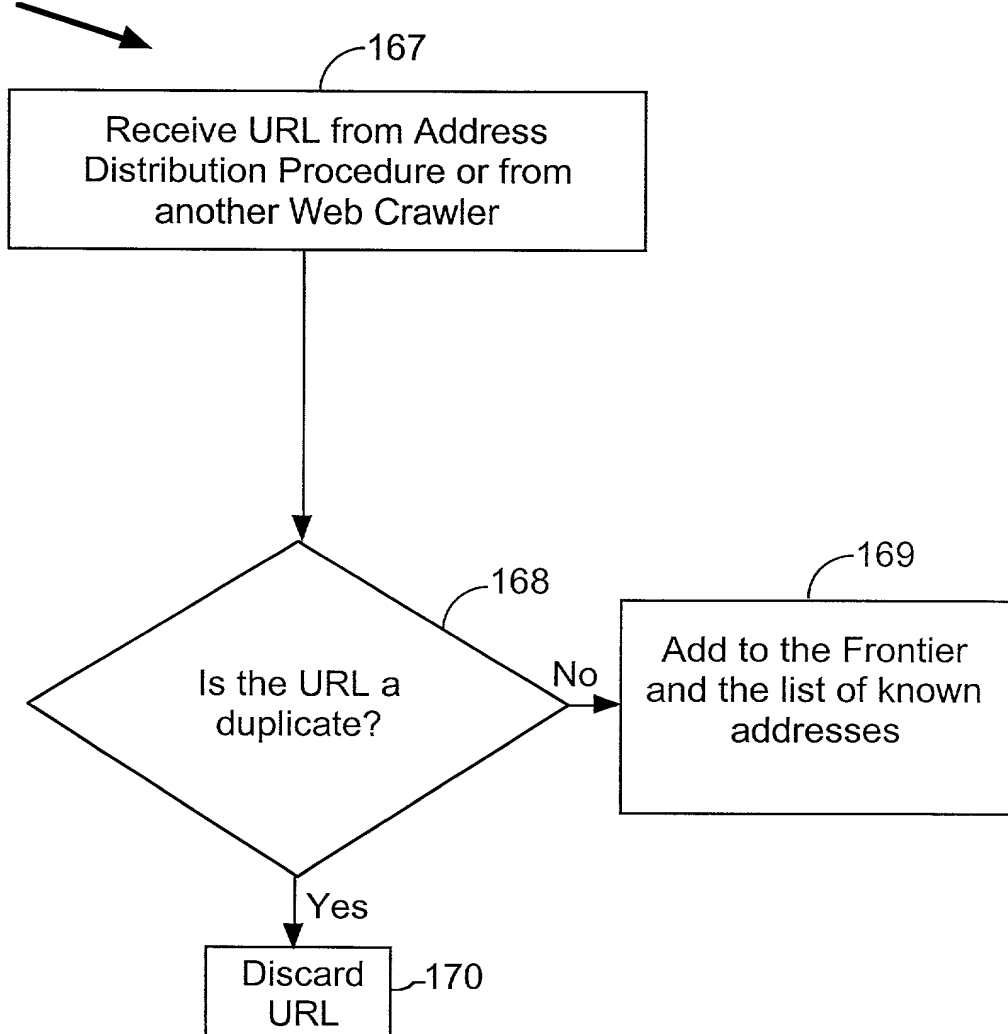


FIG. 5

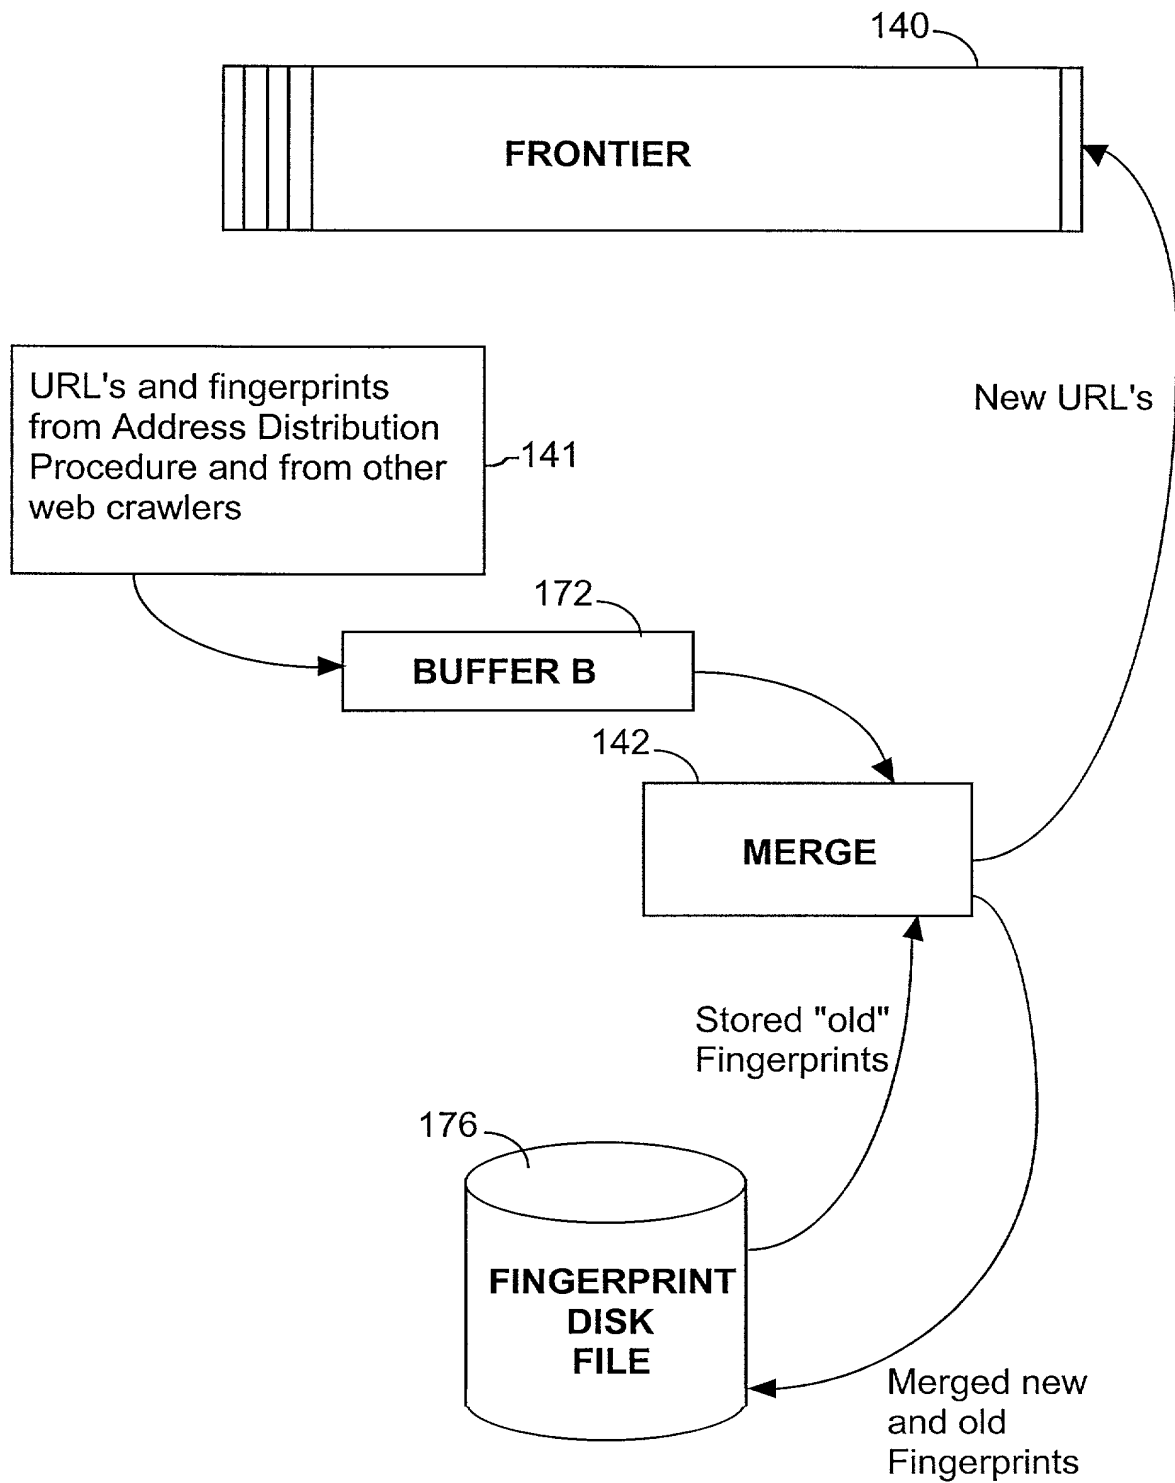


FIG. 6

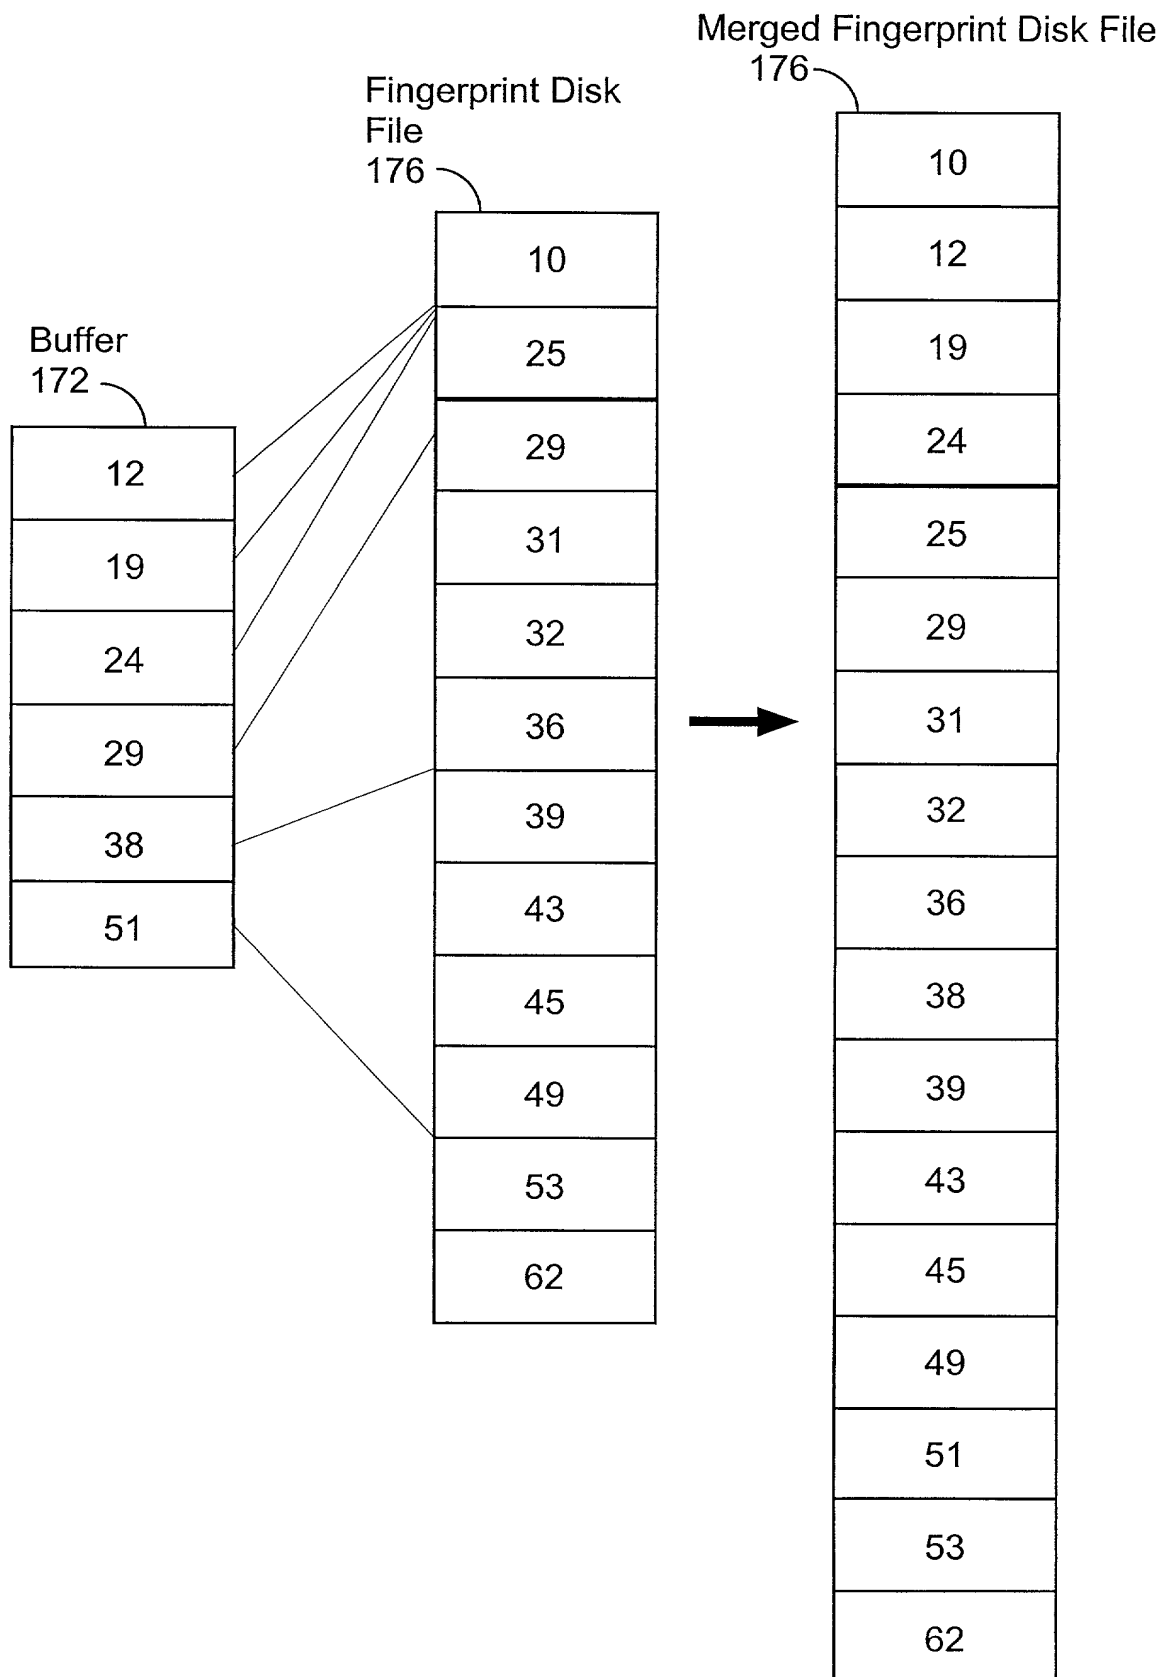


FIG. 7



FIG. 8




FIG. 9

## DECLARATION FOR PATENT APPLICATION

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled SYSTEM AND METHOD FOR DISTRIBUTED WEB CRAWLER, the specification of which

(check one) XX is attached hereto.

\_\_\_\_\_ was filed on \_\_\_\_\_ as  
Application Serial No. \_\_\_\_\_  
and was amended on \_\_\_\_\_  
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, §1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>			<u>Priority Claimed</u>	
<u>(Number)</u>	<u>(Country)</u>	<u>(Day/Month/Year Filed)</u>	<u>Yes</u>	<u>No</u>
_____	_____	_____	_____	_____
<u>(Number)</u>	<u>(Country)</u>	<u>(Day/Month/Year Filed)</u>	<u>Yes</u>	<u>No</u>
_____	_____	_____	_____	_____

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulation, §1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

<u>(Application Serial No.)</u>	<u>(Filing Date)</u>	<u>(Status)</u> (patented, pending, abandoned)
_____	_____	_____
<u>(Application Serial No.)</u>	<u>(Filing Date)</u>	<u>(Status)</u> (patented, pending, abandoned)
_____	_____	_____
<u>(Application Serial No.)</u>	<u>(Filing Date)</u>	<u>(Status)</u> (patented, pending, abandoned)
_____	_____	_____

Direct all telephone calls to **Gary S. Williams** at (650) 493-4935.  
Address all correspondence to:

PENNIE & EDMONDS LLP  
3300 Hillview Avenue  
Palo Alto, California 94304

File No. 9772-0298-999

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Title 18, United States Code, §1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of **sole or**  
**first** inventor:

MARC ALEXANDER NAJORK

Inventor's signature:

*Marc Najork*

Date:

2 November 2000

Residence:

PALO ALTO, CALIFORNIA

Citizenship:

GERMANY

Post Office Address:

67 TULIP LANE

PALO ALTO, CALIFORNIA 94303

NQU 02'00 10:36 FR PENNIE &amp; EDMONDS LLP 650 493 5556 TO 14082854065

P.02

**POWER OF ATTORNEY BY ASSIGNEE AND CERTIFICATE UNDER 37 CFR 3.73(b)**

To the Commissioner of Patents and Trademarks:

The undersigned assignee of the entire interest in application for letters patent entitled: SYSTEM AND METHOD FOR DISTRIBUTED WEB CRAWLER, and having the named inventor: MARC ALEXANDER NAJORK, hereby appoints the following attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith; said appointment to be to the exclusion of the inventor(s) and his (their) attorney(s) in accordance with the provisions of 37 C.F.R. 1.32: S. Leslie Mirock (Reg. No. 18872), Harry C. Jones, III (Reg. No. 20280), Benj A. Terzian (Reg. No. 20060), David Weild, III (Reg. No. 21094), Jonathan A. Marshall (Reg. No. 24614), Barry D. Rein (Reg. No. 22411), Stanton T. Lawrence, III (Reg. No. 25736), Charles E. McKenney (Reg. No. 22785), Philip T. Shannon (Reg. No. 24278), Francis E. Morris (Reg. No. 24815), Charles E. Miller (Reg. No. 24578), Gidon D. Stern (Reg. No. 27468), John J. Lauter, Jr. (Reg. No. 27814), Brian M. Poissant (Reg. No. 28482), Brian D. Coggio (Reg. No. 27824), Rory J. Radding (Reg. No. 28748), Stephen J. Harbulak (Reg. No. 29188), Donald J. Goodell (Reg. No. 18788), James N. Pelik (Reg. No. 25510), Thomas E. Friebe (Reg. No. 29258), Laura A. Coruzzi (Reg. No. 30742), Jennifer Gordon (Reg. No. 30753), Allan A. Fanucci (Reg. No. 30258), Geraldine F. Baldwin (Reg. No. 31232), Victor N. Balancia (Reg. No. 31231), Samuel B. Abrams (Reg. No. 30605), Steven I. Wallach (Reg. No. 35402), Marcia H. Sundeen (Reg. No. 30893), Paul J. Zegger (Reg. No. 33821), Edmond R. Bannon (Reg. No. 32110), Bruce J. Barker (Reg. No. 33291), Adriane M. Antier (Reg. No. 32605), Thomas G. Rowan (Reg. No. 34419), James G. Markey (Reg. No. 31638), Thomas D. Kohler (Reg. No. 32797), Scott D. Stimpson (Reg. No. 33807), Gary S. Williams (Reg. No. 31068), William S. Galliani (Reg. No. 33885), Ann L. Gisolfi (Reg. No. 31956), Todd A. Wagner (Reg. No. 35399), Scott B. Familant (Reg. No. 35514), Kelly D. Talcott (Reg. No. 39582), Francis D. Cerrito (Reg. No. 38100), Anthony M. Insogna (Reg. No. 35203), Brian M. Rothery (Reg. No. 35340), Brian D. Siff (Reg. No. 35678), and Alan Tenenbaum (Reg. No. 34939), all of Pennie & Edmonds LLP, whose addresses are: 1155 Avenue of the Americas, New York, New York 10036; 1667 K Street N.W., Washington, DC 20006; and, 3300 Hillview Avenue, Palo Alto, CA 94304, and each of them its attorneys, to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

In addition, the undersigned assignee also appoints: Irene Kosturakis (Reg. No. 33,724); Kath Lutsch (Reg. No. 31,851); Joseph Arrambide (Reg. No. 39,589); Sarah T. Harris (Reg. No. 35,891); Barry Blount (Reg. No. 35,068); Richard P. Lange (Reg. No. 27,286); and Theodore S. Park (Reg. No. 28,971), to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

In accordance with 37 CFR 3.73 the assignee hereby certifies that it is the assignee of the entire right, title and interest in the patent application identified above by virtue of an Assignment from the inventor(s). Said assignment recorded in the Patent Office on June 30, 2000, Reel 10840, Frames 0010-0012. The undersigned has reviewed all the documents in the chain of title of the patent application identified above and, to the best of undersigned's knowledge and belief, title is in the assignee identified below.

The undersigned hereby declares that all statements made herein of his/her own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Title 18, United States Code, 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Direct all telephone calls to GARY S. WILLIAMS at (850) 483-4836. Address all correspondence to:

PENNIE & EDMONDS LLP  
3300 HILLVIEW AVENUE  
PALO ALTO, CALIFORNIA 94304

Assignee: COMPAQ COMPUTER CORPORATION

By: Richard P. Lange

Signature: Richard P. Lange

Title: Sr. Intellectual Property Counsel

Address: COMPAQ Computer Corporation

10435 N. Tantau Blvd. 200

Sunnyvale, CA 95014-3548

Date: June 2000

File No. 9772-0298-889 (P00-3311)

CA1 -259061.1

\*\*\* TOTAL PAGE.002 \*\*\*